




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/56, 12/26		A1	(11) International Publication Number: WO 98/13975
			(43) International Publication Date: 2 April 1998 (02.04.98)
(21) International Application Number: PCT/FI97/00573		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 24 September 1997 (24.09.97)		<p>Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i> <i>In English translation (filed in Finnish).</i></p>	
(30) Priority Data:			
963797 24 September 1996 (24.09.96) FI 964268 23 October 1996 (23.10.96) FI			
(71) Applicant (for all designated States except US): NOKIA TELECOMMUNICATIONS OY [FI/FI]; Keilalahdentie 4, FIN-02150 Espoo (FI).			
(72) Inventor; and			
(75) Inventor/Applicant (for US only): MARKKI, Outi [FI/FI]; Postipuuntie 2 B 20, FIN-02600 Espoo (FI).			
(74) Agent: PATENT AGENCY COMPATENT LTD.; Teollisuuskatu 33, P.O. Box 156, FIN-00511 Helsinki (FI).			

(54) Title: PACKET ROUTING IN A TELECOMMUNICATIONS NETWORK



(57) Abstract

The invention relates to a method for packet transfer in a telecommunications network. In a node (N1) at the boundary of the network, the data packets to be transmitted are segmented into lower protocol layer data units. The data units are transferred in the network both in default channels and in dedicated channels. For efficient routing of the data packets, the node (N1) at the boundary of the network monitors packets that pertain to the same flow, and upon detecting that a given flow meets predetermined criteria, a dedicated lower protocol layer connection identifier is assigned for its use. In a node (N2) within the network, the connection identifiers of incoming data packets are monitored, and upon detection of a new connection identifier that is not associated with a default channel or a dedicated channel, an output port and a new outgoing lower protocol layer connection identifier are defined for the data unit on the basis of the destination address of that packet, and the data unit is routed to the relevant output port. Upon detection of the same lower protocol layer connection identifier as said new connection identifier in subsequently arriving data units, routing is performed in the node (N2) merely on the basis of the lower protocol layer connection identifier. Dedicated connections are discarded from use in the nodes after no traffic in a flow having a dedicated connection identifier has been detected during a given period of time.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Packet routing in a telecommunications network

Field of the Invention

The invention generally relates to packet routing carried out in a
5 packet-switched telecommunications network, particularly to the transfer of IP packets (IP, Internet Protocol) through an ATM network.

Background of the Invention

IP is the most popular of the current network layer (third layer in
10 the OSI model) protocols, mainly on account of the great popularity of the Internet. With the exponential growth of hosts connected to the Internet, the throughput of IP networks has become a bottleneck, and new ways are required to transfer IP traffic more rapidly than at present.

Figure 1 illustrates the typical structure of an IP network. In an of-
15 fice environment, personal computers PC or similar terminal equipment are connected to local area networks LAN1...LAN3, which are typically Ethernet networks. The local area networks in turn are interconnected with a backbone network (WAN, Wide Area Network), comprising routers (RT1...RT6) as nodes. All computers that are in the same local area network have the
20 same IP network address. When a data packet is sent from a computer connected to the local area network, the IP layer of the protocol stack of the sending computer checks whether the IP destination address is the same as its own IP network address. If the address is the same, no routers are necessary, but the packet is sent across the local area network to the computer
25 having that destination address. If the IP network address of the destination is different than the IP network address of the sending computer, the computer forwards the packet to a router that transfers the packet further to another network.

The transmission links between the routers can be implemented
30 with PDH or SDH technology or with the packet network technique, for instance (ATM, Frame Relay, X.25).

A router has two main functions: packet transfer and updating of routing tables. The packet transfer process in principle operates in such a way that the router first reads the network address of the destination from an
35 incoming IP packet. Thereafter it finds from its routing table the output port associated with that address and sends the packet through said port to the

next router. The packets are transferred from router to router until one of the routers finds that the destination address is the same as its own network address, in which case it sends the packet to the destination host.

With the increase in transfer rate requirements, new technologies
5 have been introduced. ATM technology is used to an ever increasing degree as a backbone network technique, as it enables high-capacity backbone connections. In such a case, interfaces have been constructed in the routing nodes towards the ATM network; packets are first reconstructed from the
10 cells arriving from the ATM network, the packets are routed, and thereafter the packets are again disassembled into ATM cells for transfer in the ATM network. A standard ATM adaptation layer (AAL) performs the disassembly and reconstruction of the IP packets. This will be described more closely hereinbelow to provide a background for the description of the invention which is to follow.

15 When a workstation in the above-described Ethernet local area network sends data to a workstation in another local area network, the data unit P1 formed by the application at the workstation is first encapsulated into a TCP packet P2, for instance, as shown in Figure 2 (provided that the protocol used in the transport layer is TCP, Transmission Control Protocol). The
20 TCP packet is thereafter encapsulated into an IP packet P3 and the IP packet further into an Ethernet frame P4 which is sent across the local area network to a router connected thereto, having an interface towards the ATM network as well. This router removes the Ethernet segment and disassembles the IP packet into ATM cells in the ATM adaptation layer. It is to be
25 noted that the encapsulation may include insertions both in front and after the packet (so-called trailer).

Figure 3 illustrates the structure of one IP packet 30 (i.e., IP datagram). The minimum size of the packet header is 20 bytes, which is divided into five four-byte "words", presented in consecutive lines in the figure. After
30 the name of each field, the length of that field in bits is shown in the figure in parentheses. The header firstly comprises a 4-bit version field 31, indicating the IP version that is used. This is followed by the length field 32 (IHL, Internet Header Length), indicating the length of the header in 4-byte words. Type field 33 indicates the type of service, and field 34 the total length of the datagram, including the header. Identification field 35 is used to identify the IP
35 packet in conjunction with the reconstruction of the packet. Flags field 36

allows the system to deduce whether a fragment of a disassembled packet is the last fragment of the original packet or not. The content of fragment offset field 37 in turn indicates in what location in the original packet the fragment belongs. Lifetime field 38 indicates the longest time that the packet may exist in the network. Each router through which the packet travels deducts from the value of this field. Protocol field 39 indicates the higher layer protocol of the data carried by the datagram (e.g. TCP). Field 40 contains a header checksum. Fields 41 and 42 are for the source and destination addresses, i.e. they indicate the address of the sending and receiving host in the form of 32-bit addresses. The address fields are followed by an option field 43 that is seldom used. The data to be transported in this field generally relates to network testing or troubleshooting; the data may for example define a given route that the datagram should travel. The field is supplemented with stuff bits when necessary, so that the number of bytes is divisible by four.

After the above-described header, the actual data in the IP packet begins. The length of the data field may vary, but its upper limit is bounded by the length of field 34, which means that the maximum length of the entire packet is 2^{16} bytes.

As stated previously, in the ATM network in accordance with Figure 1 the IP packets are transferred in the form of ATM cells. Figure 4a shows the basic structure of one cell to be transmitted in the ATM network. Each cell to be transmitted in the network comprises a 48-byte payload and a 5-byte header, but the precise structure of the header (the content of the header) is dependent on which part of the ATM network is being used in each case, as the ATM network architecture comprises a number of interfaces closely specified in the standards, and the header structure employed in the ATM cell is dependent on which interface (i.e., which part of the network) is concerned.

Figure 4b shows the cell header structure at the UNI interface (User-to-Network Interface) of the ATM network, which is the interface between the ATM terminal equipment and the ATM node. Figure 4c in turn illustrates the cell header structure at the NNI interface (Network-to-Network Interface) of the ATM network, which is the interface between two ATM nodes, either within the network or between two networks.

The routing field of the cell header consists of the virtual path identifier (VPI) and the virtual channel identifier (VCI). In the header structure

shown in Figure 4b, which is thus used at the subscriber terminal only, a total of 24 bits has been reserved for the routing field (VPI/VCI). In the header structure shown in Figure 4c, which is used in all other parts in the ATM network, 28 bits have been reserved for the routing field (VPI/VCI). As the name implies, the routing field serves as a basis for routing cells in the ATM network. Primarily in the inner parts of the network, a virtual path identifier VPI is used, which in practice often determines to what physical connection the cell is to be routed. The virtual channel identifier VCI, on the other hand, is often used for routing merely at the boundary of the network. It is to be noted, however, that it is the VPI and VCI together that unambiguously define the route for the cell.

The other fields in the ATM cell header, as defined in the specifications, are

- GFC (Generic Flow Control), a field intended for traffic supervision at the subscriber terminal, not yet precisely defined,
- PTI (Payload Type Indicator), mainly used to distinguish network management cells from subscriber information cells, but it is possible to additionally differentiate the subscriber information cells according to whether a congestion has been detected on the route or not,
- CLP (Cell Loss Priority), used to prioritize cells as regards cell loss probability (largely corresponds to the DE bit in a Frame Relay Network),
- HEC (Header Error Control), header checksum.

Of these other fields, primarily the PTI field relates to the present invention; it is possible to use this field for monitoring the limit between packets. The last bit in the PTI field (bit number two in Figures 4a...4c) indicates when a new higher layer packet (IP packet) starts. When the bit has been set to be 1, the last cell of an IP packet is involved, and in that case the next packet starts with the next non-idle cell.

In general, signals of various formats arrive at the ATM adaptation layer, and the task of the ATM adaptation layer is on the one hand to form these signals into the standard format required by the ATM network prior to their being passed on to the ATM network, and on the other hand to reconstruct these signals from the cells arriving from ATM networks prior to the signals being relayed further to the user or control interfaces. Different types of adaptation layers (AAL1...AAL5) have been standardized for different

service classes (A...D). For example, AAL types 3, 4 and 5 offer transmission services for applications in which there is no time dependency between the source and the destination.

Figure 5 illustrates the segmentation of IP packets performed by AAL 5 into ATM cells and the reconstruction of packets from ATM cells, as it illustrates the operation of the ATM adaptation layer in the exemplary network shown in Figure 1. The ATM adaptation layer is generally divided into two sublayers, which are referred to by the abbreviations SAR (Segmentation and Reassembly Sublayer) and CS (Convergence Sublayer). The last-mentioned sublayer performs encapsulation/deencapsulation of user data units (for example IP packets) and control data. The frame produced as a result of the encapsulation performed by the CS sublayer is referred to as the CS-PDU (Convergence Sublayer Protocol Data Unit). The AAL 5 encapsulation is performed in such a way that a trailer portion is added to the user data unit (e.g. an IP packet), containing for instance an error check portion (CRC). The trailer portion has a length of 8 bytes. The length of the entire CS-PDU corresponds to a multiple of 48 bytes, which is obtained by adding when appropriate a padding field PAD having a length of 0...47 bytes between the trailer portion and the payload of the packet.

The SAR sublayer segments each CS-PDU in the transmit direction into segments of 48 bytes, which are referred to as the SAR-PDU (Segmentation and Reassembly Protocol Data Unit). In the receive direction, the CS-PDUs are formed by putting the SAR-PDUs together.

The ATM layer beneath the ATM adaptation layer is responsible for adding five-byte header fields CH (Figures 4b and 4c) into the SAR-PDUs to be transmitted, thus producing ATM cells that are sent to the ATM network. The ATM network only processes the cell header, the 48-byte payload is not processed and not even read in the ATM network. In the receive direction, the ATM layer removes the headers from the cells and supplies the 48-byte payloads to the SAR sublayer for assembly.

When in the network of Figure 1 the routers send IP packets, they segment the packets into cells in the above manner and send the packets to an ATM transmission link. The router at the opposite end of the link reconstructs a packet from the cells in the above manner, makes a routing decision in the conventional manner on the basis of the IP address, and segments the packet again into cells for the next ATM link.

Normally the routing decision on the packets is made by software. When routing implemented with software is combined with the above-described packet segmentation and reconstruction, the operation of the ATM-based router network is slowed down considerably. This conventional routing method is also expensive to carry out.

To alleviate these drawbacks, a method has been developed that is termed IP switching. IP switching is based on the flow concept: a flow is a series of IP packets that are headed from (generally) the same source to the same destination. Hence, one flow (generally) comprises the IP packets whose source and destination addresses are the same. For example a TCP connection is a flow: when the TCP connection has been opened, a series of packets is sent from the source to the destination. In IP switching, router nodes identify the flows and request routers at the boundary of the network to furnish the packets of each flow with a unique flow identifier, e.g. a unique VPI/VCI identifier. When the packets pertaining to a given flow are provided with a unique VPI/VCI identifier, the routers within the network can carry out packet relay on a cell level, using a normal ATM switching. Hence, the routing need not go as far as the third layer (IP layer), but can be performed in the second layer (ATM layer).

The drawback of the last-mentioned prior art solution is, however, that it presupposes a separate control protocol wherewith the nodes within the network negotiate with the nodes in the incoming direction of traffic on assigning a dedicated VPI/VCI identifier to a given flow. Such a solution will render the network still more intricate and produce additional traffic that loads the network.

Summary of the Invention

It is an object of the invention to remove the above drawbacks and to provide a solution that will speed up packet routing without, however, requiring any new flow control protocol to be applied.

This object is achieved with the solution defined in the independent claim.

The idea of the invention is to monitor flows formed by packets in a node at the boundary of a network (receiving packets of an upper protocol layer). When it is detected that a given flow meets the criteria for a dedicated connection, a dedicated lower protocol layer connection identifier is assigned

to that flow. When the next node in the downstream direction receives a data unit having a "new" connection identifier, it immediately knows that a flow to which a dedicated channel has been assigned is concerned. The node performs normal routing on the first packet on the basis of the IP destination address, but all the following data units having the same connection identifier are forwarded on the basis of the lower protocol layer connection identifier only. Furthermore, the nodes have time-out control that discards the dedicated connections from use after no traffic pertaining to the relevant flow has been received during a predetermined period of time.

In the most preferred embodiment of the invention, the lower-layer data units are, in accordance with the foregoing, ATM cells and the higher-layer packets are IP packets, but the method can also be implemented in conjunction with other protocols or transmission methods.

On account of the solution in accordance with the invention, the throughput of nodes can be increased, obviating the need for segmenting and reconstructing packets in nodes within the network. Furthermore, this is achieved in such a way that no extra flow control protocol is needed in the network, and thus also the router nodes remain simpler and cause no additional traffic in the network.

Brief Description of the Drawings

The invention and its preferred embodiments will be described in greater detail in the following with reference to Figures 6-11g by means of examples in accordance with the accompanying drawings, in which

Figure 1 illustrates the environment in which the invention is used,

Figure 2 illustrates forming of packets prior to their being sent to the network,

Figure 3 illustrates the structure of an IP packet,

Figures 4a...4c illustrate the generic structure of an ATM cell,

Figure 5 illustrates packet segmentation into ATM cells as performed by ATM adaptation layer 5 and reconstruction of packets from ATM cells,

Figure 6 illustrates the operation of the method of the invention in three successive nodes of the network,

- Figure 7 is a flow chart illustrating the operation of the method of the invention in the first node of Figure 6,
Figure 8 is a flow chart illustrating the operation of the method of the invention in the second and third node of Figure 6,
5 Figure 9 illustrates the structure of a gateway node,
Figure 10 illustrates the structure of a node within the network, and
Figures 11a...11g show tables used in the node of Figure 10.

Detailed Description of the Invention

- 10 In the solution in accordance with the invention, segmentation of IP packets is performed in a node at the boundary of the network (either node RT1 or RT4 in Figure 1) in accordance with the above, and the ATM cells are sent to the ATM transmission link.

- Figure 6 illustrates three ATM network nodes N1...N3 which are
15 located in succession on the route for the packets and which carry out the method of the invention. In the figure, default channels between the nodes are denoted by thick pipes and dedicated channels between the nodes by thin pipes. The different steps of the method are denoted by circled figures.

- The system initially only includes default channels, and when traf-
20 fic starts, the system starts assigning dedicated channels to connections according to need. Default channels extend from a node in all directions, and packets having different source and/or destination addresses may travel in one default channel (even though they have the same VPI/VCI identifier). Hence, a default channel is a channel using ATM only as a transmission
25 path and in which routing is effected in the normal way at a higher level (IP level). For this reason, the traffic in the default channels is undesired traffic that the system attempts to eliminate. Default channels must be used, however, because in practice a large majority of all traffic is of so short a duration (possibly only a few packets are sent over the TCP connection) that there is
30 no use assigning to them a dedicated VPI/VCI identifier. It is only worth-while to establish a dedicated virtual connection for that part of the traffic which is of longer duration.

- Only IP packets arrive at the first node (N1), which segments the packets in the above-described manner into cells and relays the cells further
35 to node N2. Node N1 examines desired header fields (source and destination addresses and possibly also other fields) of the incoming IP packets in a

known manner. When node N1 detects (circled one) that a given IP connection (packets having the same source and/or destination addresses) meets certain predetermined criteria for a dedicated second-layer connection (OSI layer 2), node N1 makes a decision to assign a dedicated second-level connection (i.e., connection identifier) for that flow, which in the case of an ATM network is a new VPI/VCI identifier.

The node may make the above decision for example as a result of detecting that there is regular or heavy traffic over that IP connection (the node counts incoming packets) or upon detecting that the IP connection requires certain quality of service. The node may find, for instance, that an FTP transmission is involved which requires fast service. The above criteria may be of a wide variety of types.

When after the above decision the node N1 detects the first packet, it selects an available second-layer connection identifier (VPI/VCI identifier) on the basis of the IP header of that packet.

For example, when the node has detected at the tenth packet that the connection requires a dedicated second-layer connection, the node N1 starts using the identifier appertaining to the dedicated second-layer connection in the cells it disassembles (circled two), starting with said packet. The first nine packets have been forwarded in the default channel (in the form of cells). The cells in each packet are sent in succession in the default channel in such a way that the discrete packets do not overlap.

After this, the next node (N2) in the downstream knows (circled three), having received a cell carrying a VPI/VCI identifier that is "new" to the node (not used by any one of the flows), that a flow is concerned to which a dedicated second-layer connection identifier has been assigned in the node in the upstream direction. From this cell the node reads the IP destination address that is mapped, as can be seen from Figure 3, in segmentation in such a way that it always has a fixed location in the payload bytes 13-17 of the cell. On the basis of the destination address that has been read, the node finds the identifier of the output port associated with that address from the routing table and transmits the cell to that output port. Thereafter all cells having the same VPI/VCI identifier are transmitted to the same output port. The node only needs to check the first IP address, and it performs normal routing (at the IP level) on the basis thereof. After this, it can perform switching at the ATM level directly on the basis of the VPI/VCI identifiers and

transmit the cells again with a new VPI/VCI identifier to the next node (circled four).

Figure 7 presents a flow chart illustrating the operation of a node at the boundary of the network (gateway node N1). In the figure, references
5 T_n in parentheses (n=1...5) are shown in the different steps; these refer to tables used by the node, which are set forth below in Figure 9.

First, the node receives an IP packet, reads from its header at least the source and destination addresses (step 71), and updates the packet counter (step 72, Table T1) for that address pair (or the destination
10 address only).

After this, the node examines whether a dedicated connection has been set for that address pair (step 73). This is done by looking up in the table whether a dedicated channel pertaining to that address pair is set on (Table T1), in other words, whether a dedicated second-layer connection
15 identifier has been assigned to that flow. If this is the case, the routing tag associated with the connection is sought from the routing table (table T2); the tag is the internal identifier of the ATM switching matrix on the basis of which the ATM switching fabric performs switching to the correct output port (step 74).

20 If no dedicated connection has been assigned yet, the next step is to examine whether the criteria for a dedicated connection are met for that flow (step 75), in other words, whether the traffic volume has exceeded a given value during a measurement period or whether the quality requirements for the traffic in the flow are of a given type (i.e. require a high
25 throughput probability), for example. The node may utilize for instance the connection type field 33 (Figure 3), indicative of the quality of service in terms of delay, throughput, and reliability. If these criteria are not met, the routing tag of the default channel for the relevant IP destination address is sought (step 74, table T2), and routing is performed accordingly.

30 If it is found in step 75 that the criteria for a dedicated connection are met, the output port associated with the relevant address pair or destination address (Table T2) and an available routing tag corresponding to the output port (table T3), i.e. a routing tag that has not been assigned yet and that corresponds to the IP address, are sought. Thereafter a dedicated
35 channel corresponding to the IP address is set on (table T1) and the routing tag for that connection is updated in table T2. At the output port, an available

VPI/VCI is sought for the new routing tag (table T5) and updated in the routing tag translation table (table T4).

5 After this, the method can proceed to step 78 in which the IP packet is segmented into cells and the routing tag that was retrieved above is attached to the cells. Thereafter the cells are transmitted to the output port associated with the routing tag (step 79). Prior to sending of the cells to the next link, the routing tag is replaced with the corresponding VPI/VCI identifier (table T4).

Thereafter the method proceeds to receiving the next IP packet.

10 In addition, the node has a separate aging logic that discards dedicated connections and connection data relating to them from use after no traffic pertaining to that particular flow has been received during a predetermined period of time. It is preferable to use a similar mechanism in all nodes, and hence the deletion of a dedicated connection is initiated from the starting node. If the node has already discarded the dedicated connection and there is traffic in the flow, the routing decision is made as described above similarly as for the "first" packet.

Figure 8 illustrates the operation of a node within a network (for example nodes N2 and N3). In the figure, references Sn in parentheses (n=1...7) are shown in the different steps; these refer to tables used by the node, which are set forth below in Figures 10 and 11a...11g.

Initially the node receives a cell (step 80) and reads its header (step 81). Thereafter the traffic counter pertaining to the VPI/VCI identifier is updated (step 82, table S1). As a next step, the node examines whether a default channel is concerned, i.e. whether the incoming VPI/VCI identifier is marked as the default channel (step 83, table S1).

If a default channel is not involved, the next step is to examine whether the flow is a flow for which a dedicated channel has been set on (step 84, table S1), i.e. whether the incoming VPI/VCI has been marked as a dedicated channel. If this is the case, the routing tag associated with the incoming VPI/VCI identifier is taken (step 86) directly from table S3.

If, on the other hand, no dedicated channel has been assigned to the flow yet, the method proceeds to step 88 in which the IP destination address is read from the cell; this address is mapped, as is seen from Figure 3, in segmentation in such a way that it always has a fixed location in the payload bytes 13-17 of the cell. On the basis of the destination address that

has been read, the output port associated with this address (table S4) and a corresponding available routing tag (table S5) are sought from the routing table. In the output port, an available VPI/VCI identifier is sought for the routing tag (table S7); this identifier is set in the routing tag translation table of the output port (table S6). At the same time, the connection is marked as
5 dedicated and the data is updated in the tables (tables S1, S3, S4).

If it is found in step 83 that a default channel is concerned, the traffic counter pertaining to the IP address pair or IP destination address is updated (step 85, table S2) and it is examined on the basis of the counter
10 value whether the criteria for a dedicated connection are met for that flow (step 87). If the criteria for a dedicated connection are not met, the cell is routed in the normal way, i.e. a route (output port) for that particular IP destination address and a corresponding routing tag are sought (step 89, table S4). If, on the other hand, the criteria for a dedicated connection are met, the
15 procedure is the same as in connection with setting up a new dedicated channel (step 88).

After the steps described above, the node has found the correct routing tag and the cell can be routed to its correct output port (step 90). Prior to sending of the cells to the next link, the routing tag is removed and
20 the header is provided with an outgoing VPI/VCI identifier (table S6).

As is apparent from the flow chart presented above, it is also possible to set up dedicated channels in a node within a network for example on the basis of the destination address, as the traffic may be multiplexed from several sources or parallel links to the same destination. If multiplexing is
25 employed on dedicated channels, this must be done preserving the order of the packets.

Figure 9 illustrates a possible implementation of a gateway node in the form of a functional block diagram.

For temporary storage of incoming packets, a buffer 91 is provided at the input of the node. A measuring and control block 92 reads the
30 header of the packet and finds, by means of the source/destination address pair, from its table (T1) whether the dedicated channel parameter has been set on for said address pair.

Retrieval block 93 receives the IP address from the measuring and control block and, on the basis thereof, seeks from routing tables (T2) a
35 routing tag that the retrieval block supplies to a segmentation block 95 in



which the segmentation of cells and attachment of routing tags to segmented cells is carried out.

Furthermore, bookkeeping units 98 that keep track of available routing tags associated with each input and output port (table T3) are provided on the input side of the node.

The cell is sent with the routing tag from the segmentation unit to an ATM switch 96, which connects the cell on the basis of the routing tag to the correct output port 97. Hence, the routing tag is an identifier within the switch, which is not sent forward in the network.

The measuring and control block measures the traffic in each flow and deletes the parameters of a dedicated channel, such as the routing tag and the VPI/VCI, when it has not detected any traffic pertaining to that flow during a given period of time.

Figure 10 is a schematic representation of the structure of a node within a network. The blocks in the node are the same as above, only the tables used by the node are different from those of a gateway node. For this reason, the different blocks are denoted by the same reference numerals as the corresponding blocks in Figure 9. The tables S1...S7 used in the different blocks have been separately set forth in Figures 11a...11g.

Measuring and control block 92 maintains for each VPI/VCI pair a counter and information on whether the relevant channel is a default channel or a dedicated channel (table S1, Figure 11a). Furthermore, in order for it to be able to transfer traffic of a default channel to a dedicated channel, it maintains a counter for the IP address pair or IP destination address (table S2, Figure 11b). Within the network, the traffic of the source/destination pair may be multiplexed for example from parallel links, or traffic may be multiplexed from several sources to a given destination address for which a dedicated channel is set up.

The routing tags of the dedicated channels, corresponding to the VPIs/VCIs, are provided in a separate table (S3, Figure 11c), wherefrom they are available for use by the routing block. Furthermore, a routing table is provided for all connections, in which an IP address is associated with an output port and a routing tag (table S4, Figure 11d). This table is needed for all routed packets in the default channel and for setting up a new dedicated channel.

Bookkeeping unit 98 keeps track of available routing tags associated with each input port and output port (table S5, Figure 11e).

5 The ATM output port has a translation table in which the routing tag of an incoming cell is replaced with a VPI/VCI for the outbound cell (table S6, Figure 11f). Furthermore, the output port keeps track of available VPI/VCI identifiers (table S7, Figure 11g).

10 Even though the invention has been explained in the above with reference to examples in accordance with the accompanying drawings, it is to be understood that the invention is not restricted thereto, but it may be modified within the scope of the inventive idea set forth in the appended claims. For example, a flow may comprise, in accordance with the above, traffic between the same source/destination pair or traffic bound to the same destination.

Claims

1. A method for packet transfer in a packet-switched telecommunications network in which packets are transferred from one node to another in a network on the basis of an address contained in each packet, in accordance with which method

- in a node (N1) located at the boundary of the network, data packets (30) to be transferred are segmented into lower protocol layer data units (50),

- the data units are transferred in the network both in default channels and in dedicated channels, data units of packets that have different destination addresses being transferred in a discrete default channel and data packets that pertain to the same flow being transferred in a discrete dedicated channel,

characterized in that

- the node (N1) at the boundary of the network monitors packets pertaining to the same flow, and upon detection of the fact that a given flow meets predetermined criteria, a dedicated lower protocol layer connection identifier is assigned for the use of that flow,

- a node (N2) within the network monitors the connection identifiers of incoming data packets, and upon detection of a new connection identifier that is not associated with a default channel or a dedicated channel, an output port and a new outgoing lower protocol layer connection identifier are assigned to the data unit on the basis of the destination address of that packet, and the data unit is routed to the relevant output port,

- upon detection of the same lower protocol layer connection identifier as said new connection identifier in subsequently arriving data units, routing is performed in the node (N2) merely on the basis of the lower protocol layer connection identifier, and

- dedicated connections are discarded from use in the nodes after no traffic in a flow having a dedicated connection identifier has been detected during a given period of time.


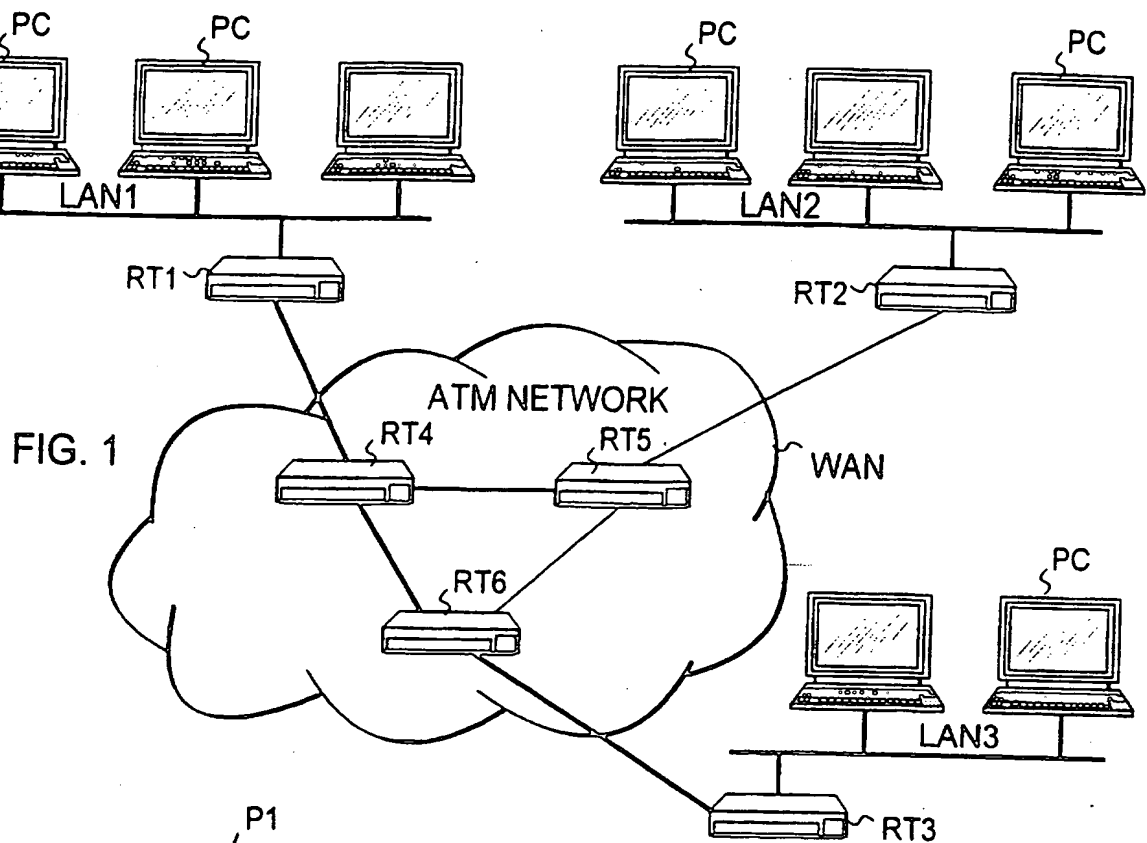
2. A method as claimed in claim 1, characterized in that similar time-out control is used in each node.

3. A method as claimed in claim 1, characterized in that also the node (N2) within the network monitors traffic in the default channels and assigns dedicated lower protocol layer connection identifiers to the con-

nections upon detecting that a given flow of a default channel meets predetermined criteria.

4. A method as claimed in claim 1, characterized in that the quantity of traffic received during a given measurement period is used as said predetermined criteria.
5. A method as claimed in claim 1, characterized in that information contained in the data packets on the quality of service required by them is used as said predetermined criteria.

1/7



2/7

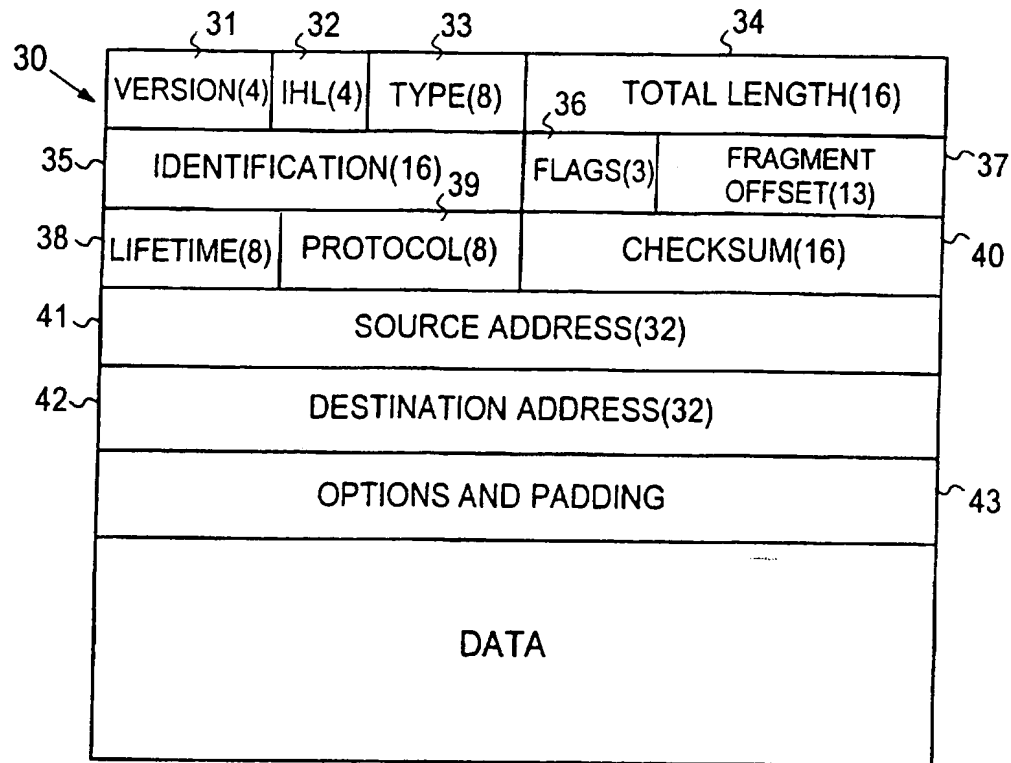


FIG. 3

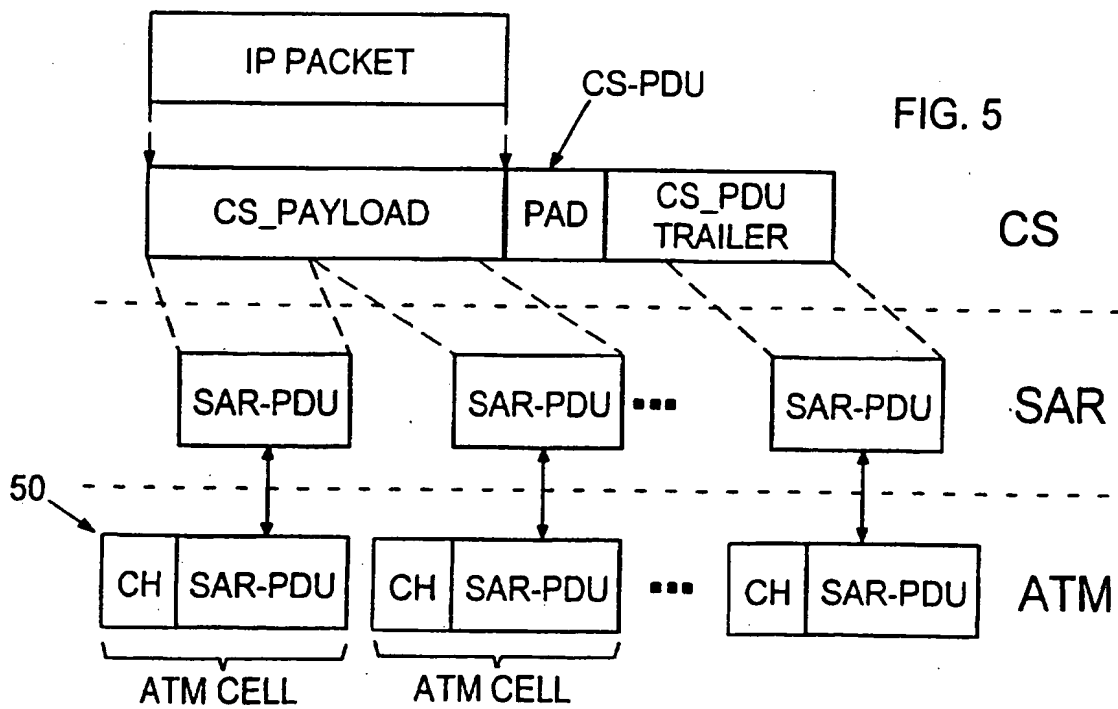


FIG. 5

3/7

BITS: 8 7 6 5 4 3 2 1




FIG. 4a

8	7	6	5	4	3	2	1	BIT BYTE
GFC				VPI				1
VPI				VCI				2
VCI								3
VCI				PTI		CLP		4
HEC								5

FIG. 4b

8	7	6	5	4	3	2	1	BIT BYTE
VPI								1
VPI				VCI				2
VCI								3
VCI				PTI		CLP		4
HEC								5

FIG. 4c

4/7

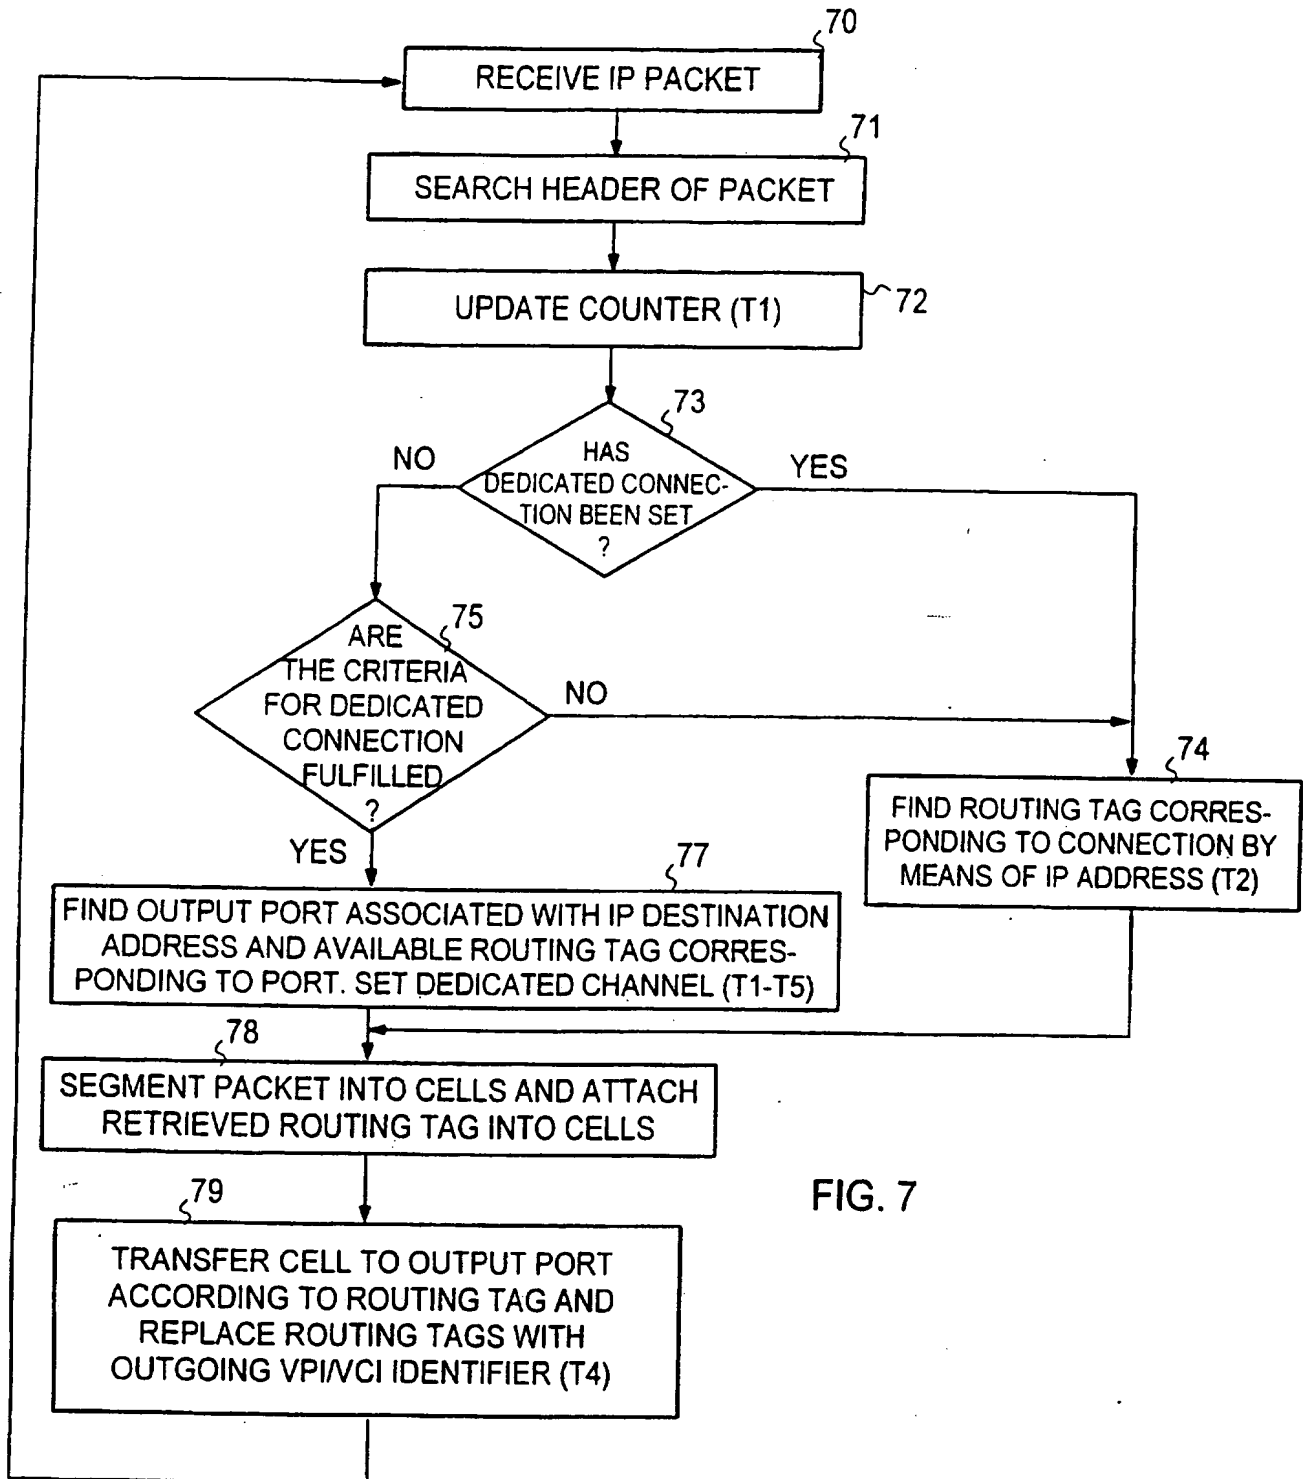
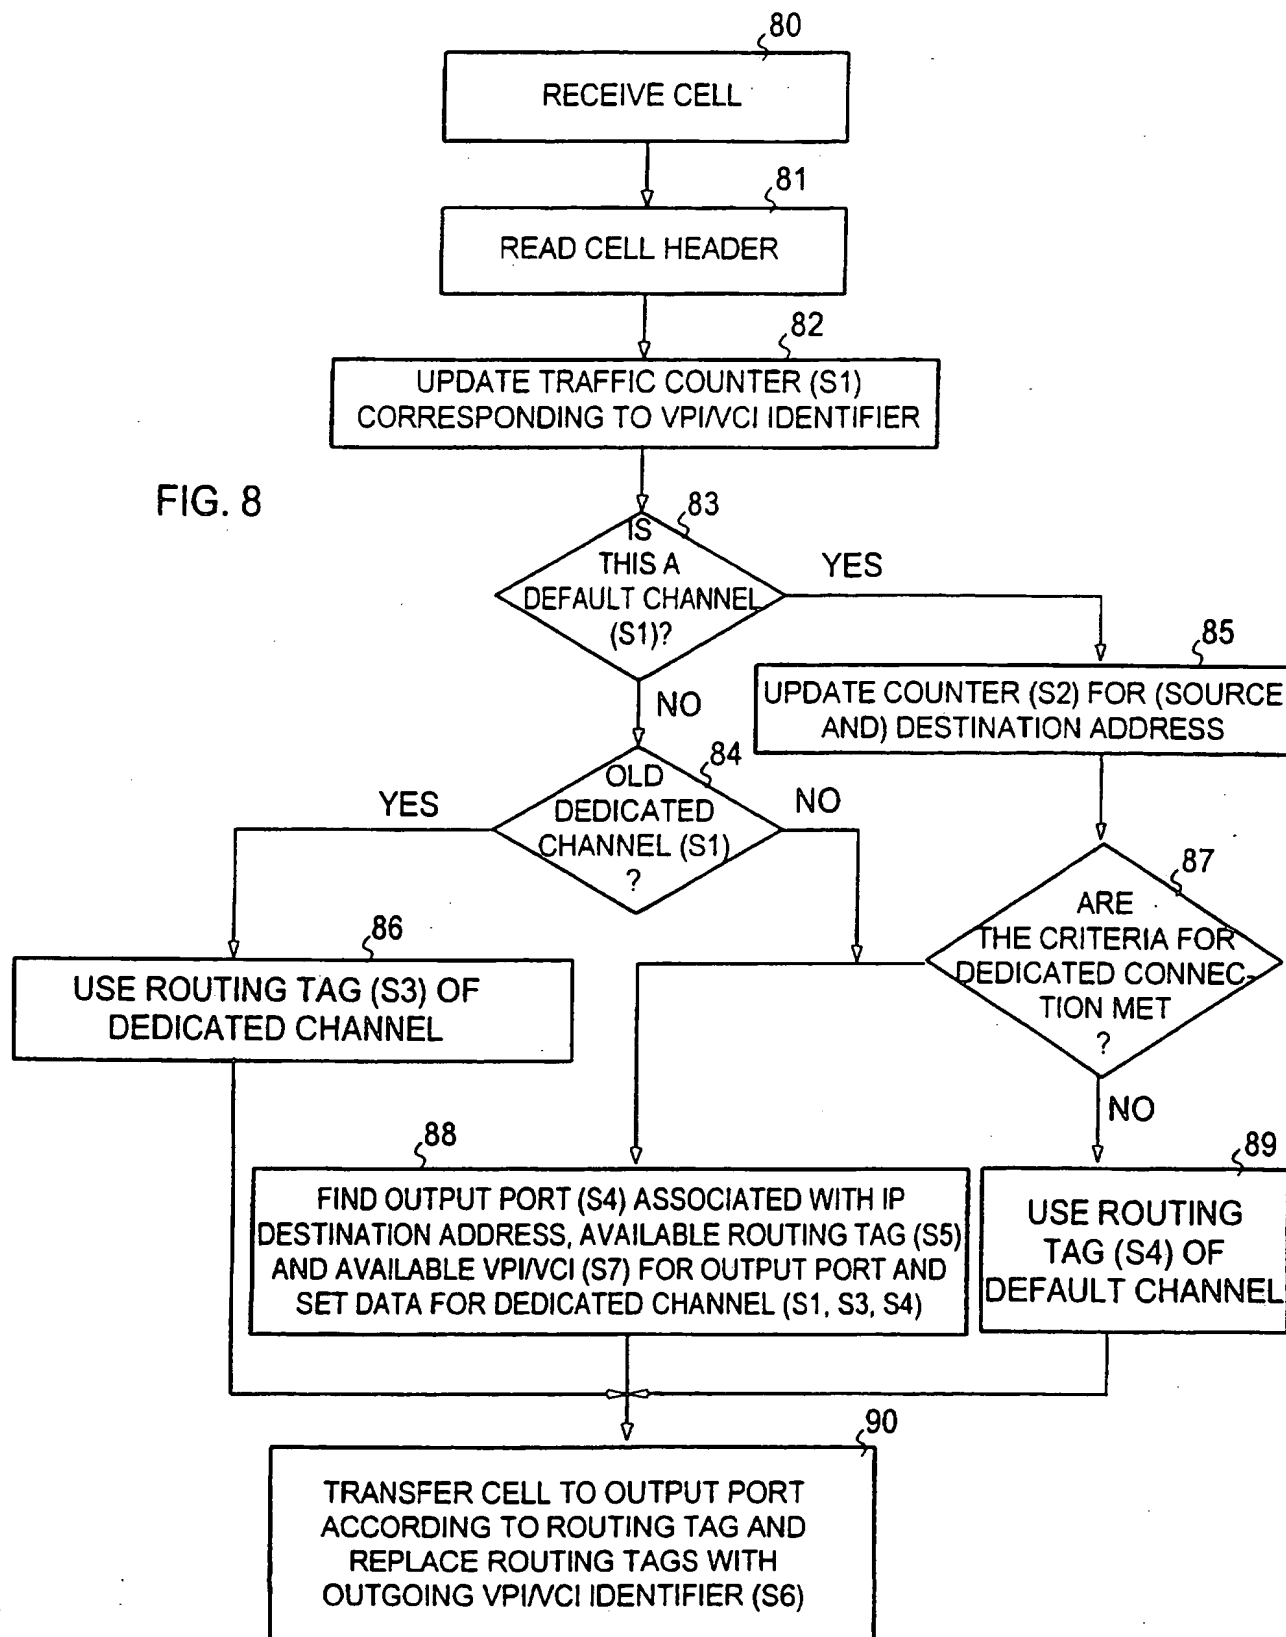


FIG. 7

5/7



6/7




FIG. 6

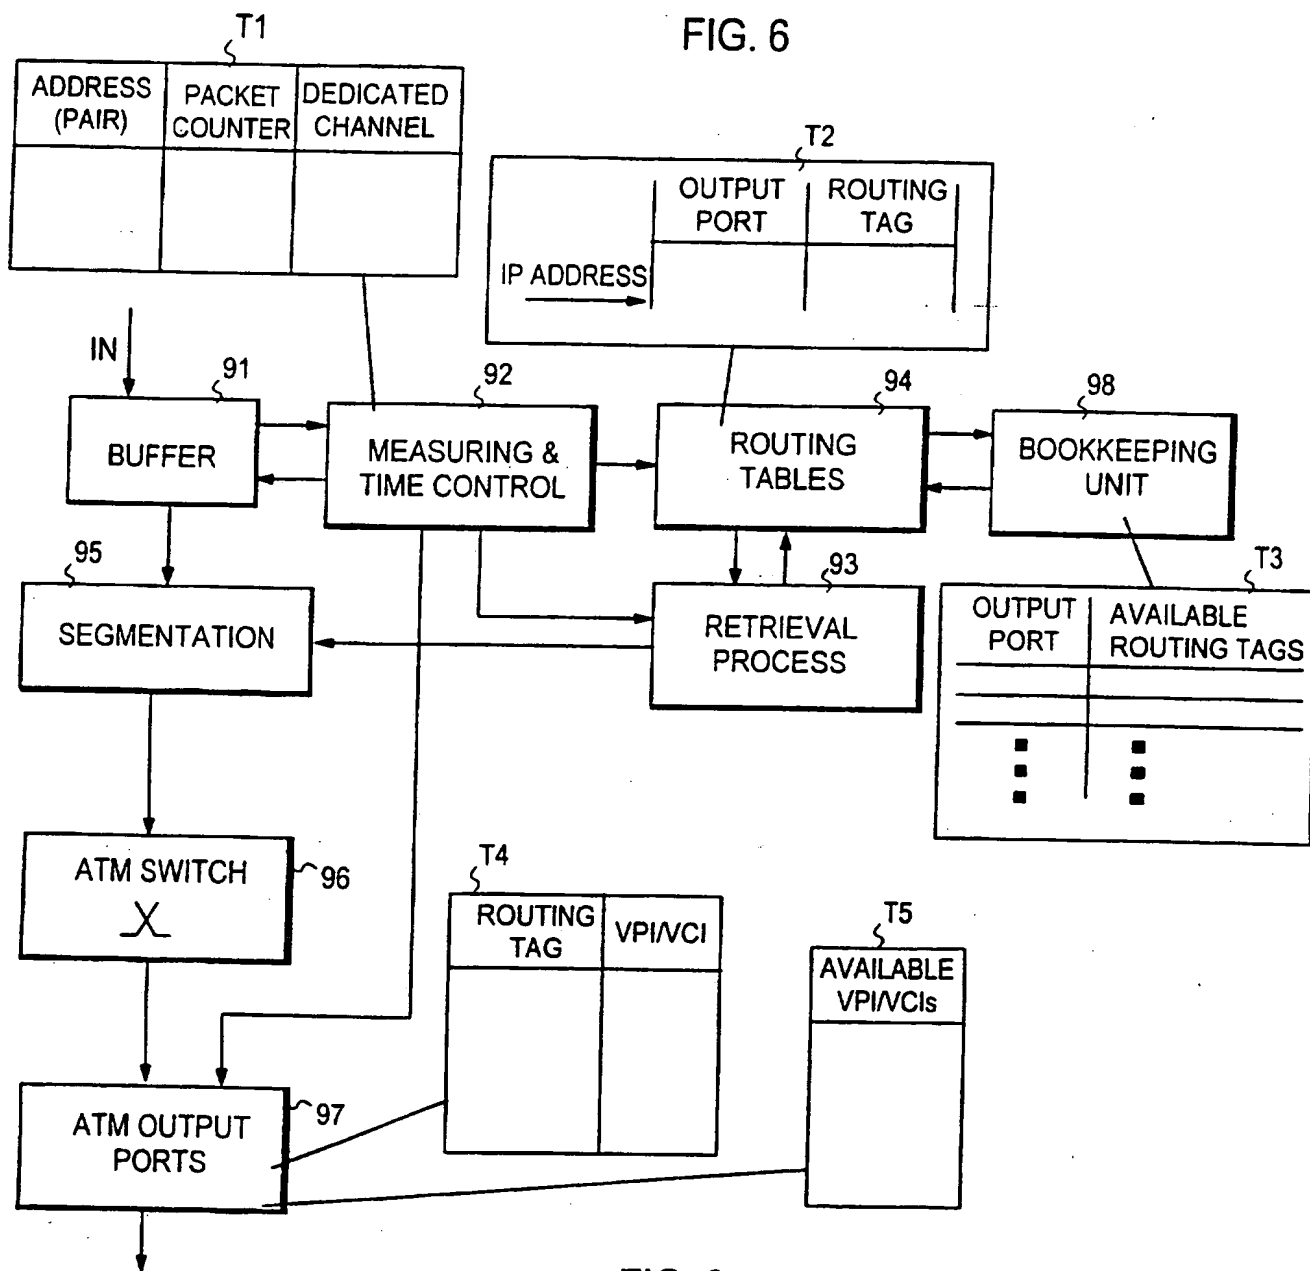


FIG. 9

7/7




FIG. 10

FIG. 11a

VPI/VC1	COUNTER	DEFAULT CHANNEL	DEDICATED CHANNEL
□	□	□	□
□	□	□	□
□	□	□	□

FIG. 11b

IP DESTINATION ADDRESS	COUNTER
□	□
□	□
□	□

IP SOURCE ADDRESS →

FIG. 11c

ROUTING TAG
□
□
□

VPI/VC1 →

FIG. 11d

OUTPUT PORT	ROUTING TAG
□	□
□	□
□	□

IP ADDRESS →

FIG. 11e

OUTPUT PORT	AVAILABLE ROUTING TAGS
□	□
□	□
□	□

FIG. 11f

ROUTING TAG	VPI/VC1
□	□
□	□
□	□

FIG. 11g

AVAILABLE VPI/VC1s
□
□
□

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 97/00573

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: H04L 12/56, H04L 12/26

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0597487 A2 (NEC CORPORATION), 18 May 1994 (18.05.94), see the whole document --	1-5
A	EP 0740442 A2 (SUN MICROSYSTEMS INC.), 30 October 1996 (30.10.96), page 2, line 37 - page 3, line 13 --	1-5
A	US 5463621 A (HIROSHI SUZUKI), 31 October 1995 (31.10.95), column 1, line 44 - column 2, line 29 --	1-5
P,A	EP 0756435 A2 (FUJITSU LIMITED), 29 January 1997 (29.01.97), page 2, line 51 - page 3, line 9; page 6, line 27 - page 7, line 15 --	1-5

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "B" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

25 February 1998

Date of mailing of the international search report

26 -02- 1998

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Friedrich Kühn
Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT
Information on patent family members

03/02/98

International application No.

PCT/FI 97/00573

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0597487 A2	18/05/94	AU 664494 B AU 5060193 A JP 6152634 A JP 7107990 B US 5452296 A	16/11/95 26/05/94 31/05/94 15/11/95 19/09/95
EP 0740442 A2	30/10/96	JP 9130401 A US 5579480 A	16/05/97 26/11/96
US 5463621 A	31/10/95	CA 2124424 A JP 2518515 B JP 6338901 A	28/11/94 24/07/96 06/12/94
EP 0756435 A2	29/01/97	JP 9046344 A JP 9064891 A JP 9162897 A	14/02/97 07/03/97 20/06/97

THIS PAGE BLANK (USPTO)